

compare-mt: A Tool for Holistic Comparison of Language Generation Systems

Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, Xinyi Wang
Language Technologies Institute, Carnegie Mellon University
gneubig@cs.cmu.edu

Abstract

In this paper, we describe `compare-mt`, a tool for *holistic* analysis and comparison of the results of systems for language generation tasks such as machine translation. The main goal of the tool is to give the user a high-level and coherent view of the salient differences between systems that can then be used to guide further analysis or system improvement. It implements a number of tools to do so, such as analysis of accuracy of generation of particular types of words, bucketed histograms of sentence accuracies or counts based on salient characteristics, and extraction of characteristic n -grams for each system. It also has a number of advanced features such as use of linguistic labels, source side data, or comparison of log likelihoods for probabilistic models, and also aims to be easily extensible by users to new types of analysis. `compare-mt` is a pure-Python open source package,¹ that has already proven useful to generate analyses that have been used in our published papers.

1 Introduction

Tasks involving the generation of natural language are ubiquitous in NLP, including machine translation (MT; Koehn (2010)), language generation from structured data (Reiter and Dale, 2000), summarization (Mani, 1999), dialog response generation (Oh and Rudnicky, 2000), image captioning (Mitchell et al., 2012). Unlike tasks that involve prediction of a single label such as text classification, natural language texts are nuanced, and there are not clear yes/no distinctions about whether outputs are correct or not. Evaluation measures such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Denkowski and Lavie, 2011), and many others attempt to give an

¹Code <http://github.com/neulab/compare-mt> and video demo <https://youtu.be/K-MNPOGKnDQ> are available.

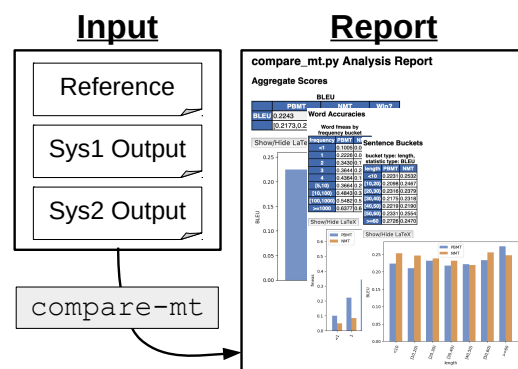


Figure 1: Workflow of using `compare-mt` for analysis of two systems

overall idea of system performance, and technical research often attempts to improve accuracy according to these metrics.

However, as useful as these metrics are, they are often opaque: if we see, for example, that an MT model has achieved a gain in one BLEU point, this does not tell us what characteristics of the output have changed. Without fine-grained analysis, readers of research papers, or even the writers themselves can be left scratching their heads asking “*what exactly is the source of the gains in accuracy that we’re seeing?*”

Unfortunately, this analysis can be time-consuming and difficult. Manual inspection of individual examples can be informative, but finding salient patterns for unusual phenomena requires perusing a large number of examples. There is also a risk that confirmation bias will simply affirm pre-existing assumptions. If a developer has some hypothesis about specifically what phenomena their method should be helping with, they can develop scripts to automatically test these assumptions. However, this requires deep intuitions with respect to what changes to expect in advance, which cannot be taken for granted in beginning researchers

or others not intimately familiar with the task at hand. In addition, creation of special-purpose one-off analysis scripts is time-consuming.

In this paper, we present `compare-mt`, a tool for *holistic comparison and analysis* of the results of language generation systems. The main use case of `compare-mt`, illustrated in 1, is that once a developer obtains multiple system outputs (e.g. from a baseline system and improved system), they feed these outputs along with a reference output into `compare-mt`, which extracts aggregate statistics comparing various aspects of these outputs. The developer can then quickly browse through this holistic report and note salient differences between the systems, which will then guide fine-grained analysis of specific examples that elucidate exactly what is changing between the two systems.

Examples of the aggregate statistics generated by `compare-mt` are shown in §2, along with description of how these lead to discovery of salient differences between systems. These statistics include word-level accuracies for words of different types, sentence-level accuracies or counts for sentences of different types, and salient n -grams or sentences where one system does better than the other. §4 demonstrates `compare-mt`’s practical applicability by showing some case studies where has *already* been used for analysis in our previously published work. Appendix A further details more advanced functionality of `compare-mt` that can make use of specific labels, perform analysis over source side text through alignments, and allow simple extension to new types of analysis. The methodology in `compare-mt` is inspired by several previous works on *automatic error analysis* (Popović and Ney, 2011), and we perform an extensive survey of the literature, note how many of the methods proposed in previous work can be easily realized by using functionality in `compare-mt`, and detail the differences with other existing toolkits in Appendix B.

2 Basic Analysis using `compare-mt`

Using `compare-mt` with the default settings is as simple as typing

```
compare-mt ref sys1 sys2
```

where `ref` is a manually curated reference file, and `sys1` and `sys2` are the outputs of two systems that we would like to compare. These analy-

	PBMT	NMT	Win?
BLEU	22.43 [21.76,23.19]	24.03 [23.33,24.65]	$s2 > s1$ $p < 0.001$
RIBES	80.00 [79.39,80.64]	80.00 [79.44,80.92]	- $p = 0.44$
Length	94.79 [94.10,95.49]	93.82 [92.90,94.85]	$s1 > s2$ $p < 0.001$

Table 1: Aggregate score analysis with scores, confidence intervals, and pairwise significance tests.

sis results can be written to the terminal in text format, but can also be written to a formatted HTML file with charts and LaTeX tables that can be directly used in papers or reports.²

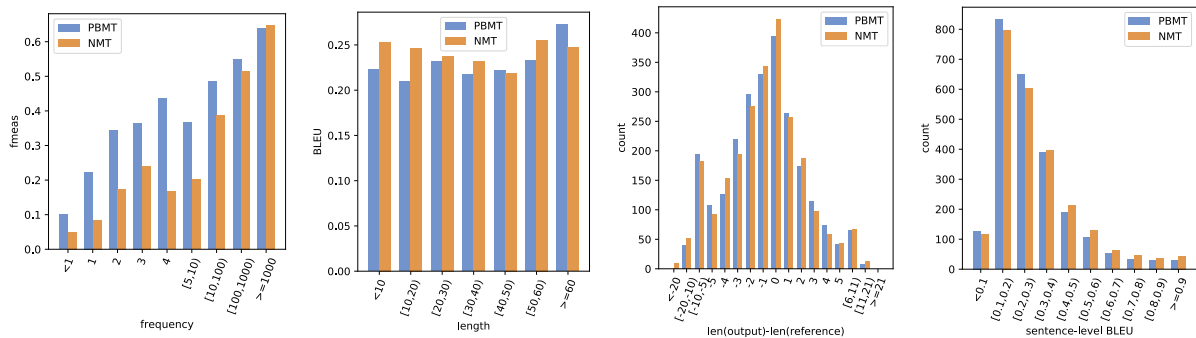
In this section, we demonstrate the types of analysis that are provided by this standard usage of `compare-mt`. Specifically, we use the example of comparing phrase-based (Koehn et al., 2003) and neural (Bahdanau et al., 2015) Slovak-English machine translation systems from Neubig and Hu (2018).

Aggregate Score Analysis The first variety of analysis is not unique to `compare-mt`, answering the standard question posed by most research papers: “*given two systems, which one has better accuracy overall?*” It can calculate scores according to standard BLEU (Papineni et al., 2002), as well as other measures such as output-to-reference length ratio (which can discover systematic biases towards generating too-long or too-short sentences) or alternative evaluation metrics such as chrF (Popović, 2015) and RIBES (Isozaki et al., 2010). `compare-mt` also has an extensible `Scorer` class, which will be used to expand the metrics supported by `compare-mt` in the future, and can be used by users to implement their own metrics as well. Confidence intervals and significance of differences in these scores can be measured using bootstrap resampling (Koehn, 2004).

Fig. 1 shows the concrete results of this analysis on our PBMT and NMT systems. From the results we can see that the NMT achieves higher BLEU but shorter sentence length, while there is no significant difference in RIBES.

Bucketed Analysis A second, and more nuanced, variety of analysis supported by

²In fact, all of the figures and tables in this paper (with the exception of Fig. 1) were generated by `compare-mt`, and only slightly modified for formatting. An example of the command used to do so is shown in the Appendix.



(a) Word F-measure bucketed by frequency. (b) BLEU bucketed by sentence length. (c) Counts of sentences by length difference. (d) Counts of sentences by sentence-level BLEU.

Figure 2: Examples of bucketed analysis

`compare-mt` is *bucketed* analysis, which assigns words or sentences to buckets, and calculates salient statistics over these buckets.

Specifically, bucketed word accuracy analysis attempts to answer the question “*which types of words can each system generate better than the other?*” by calculating word accuracy by bucket. One example of this, shown in Fig. 2a, is measurement of word accuracy bucketed by frequency in the training corpus. By default this “accuracy” is defined as f-measure of system outputs with respect to the reference, which gives a good overall picture of how well the system is doing, but it is also possible to separately measure precision or recall, which can demonstrate how much a system over- or under-produces words of a specific type as well. From the results in the example, we can see that both PBMT and NMT systems do more poorly on rare words, but the PBMT system tends to be more robust to low-frequency words while the NMT system does a bit better on very high-frequency words.

A similar analysis can be done on the sentence level, attempting to answer questions of “*on what types of sentences can one system perform better than the other?*” In this analysis we define the “bucket type”, which determines how we split sentences into bucket, and the “statistic” that we calculate for each of these buckets. For example, `compare-mt` calculates three types of analysis by default:

- **bucket=length, statistic=score:** This calculates the BLEU score by reference sentence length, indicating whether a system does better or worse at shorter or longer sentences. From the Fig. 2b, we can see that the PBMT

system does better at very long sentences, while the NMT system does better at very short sentences.

- **bucket=lengthdiff, statistic=count:** This outputs a histogram of the number of sentences that have a particular length difference from the reference output. A distribution peaked around 0 indicates that a system generally matches the output length, while a flatter distribution indicates a system has trouble generating sentences of the correct length Fig. 2c indicates that while PBMT rarely generates extremely short sentences, NMT has a tendency to do so in some cases.
- **bucket=score, statistic=count:** This outputs a histogram of the number of sentences receiving a particular score (e.g. sentence-level BLEU score). This shows how many sentences of a particular accuracy each system outputs. Fig. 2d, we can see that the PBMT system has slightly more sentences with low scores.

These are just three examples of the many different types of sentence-level analysis that are possible with difference settings of the bucket and statistic types.

N-gram Difference Analysis The holistic analysis above is quite useful when word or sentence buckets can uncover salient accuracy differences between the systems. However, it is also common that we may not be able to predict *a-priori* what kinds of differences we might expect between two systems. As a method for more fine-grained analysis, `compare-mt` implements a method that looks at differences in the n -grams produced by

<i>n</i> -gram	m_1	m_2	s
phantom	34	1	0.945
Amy	9	0	0.909
, who	8	0	0.900
my mother	7	0	0.889
else happened	5	0	0.857
going to show you	0	6	0.125
going to show	0	6	0.125
hemisphere	0	5	0.143
Is	0	5	0.143
'm going to show	0	5	0.143

Table 2: Examples discovered by n -gram analysis

each system, and tries to find n -grams that each system is better at producing than the other (Akabe et al., 2014). Specifically, it counts the number of times each system matches each ngram \mathbf{x} , defined as $m_1(\mathbf{x})$ and $m_2(\mathbf{x})$ respectively, and calculates a smoothed probability of an n -gram match coming from one system or another:

$$p(\mathbf{x}) = \frac{m_1(\mathbf{x}) + \alpha}{m_1(\mathbf{x}) + m_2(\mathbf{x}) + 2\alpha}. \quad (1)$$

Intuitively, n -grams where the first system excels will have a high value (close to 1), and when the second excels the value will be low (close to 0). If smoothing coefficient α is set high, the system will prefer frequent n -grams with robust statistics, and when α is low, the system will prefer highly characteristic n -grams with a high ratio of matches in one system compared to the other.

An example of n -grams discovered with this analysis is shown in Tab. 2. From this, we can then explore the references and outputs of each system, and figure out what phenomena resulted in these differences in n -gram accuracy. For example, further analysis showed that the relatively high accuracy of “hemisphere” for the NMT system was due to the propensity of the PBMT system to output the mis-spelling “hemispher,” which it picked up from a mistaken alignment. This may indicate the necessity to improve alignments for word stems, a problem that could not have easily been discovered from the bucketed analysis in the previous section.

Sentence Example Analysis Finally, `compare-mt` makes it possible to analyze and compare individual sentence examples based on statistics, or differences of statistics. Specifically, we can calculate a measure of accuracy of each sentence (e.g. sentence-level BLEU score), sort the sentences in the test set according to the difference in this measure, then display the

Ref/Sys	BLEU	Text
Ref	-	Beth Israel 's in Boston .
PBMT	1.00	Beth Israel 's in Boston .
NMT	0.41	Beat Isaiill is in Boston .
Ref	-	And what I 'm talking about is this .
PBMT	0.35	And that 's what I 'm saying is this .
NMT	1.00	And what I 'm talking about is this .

Table 3: Sentence-by-sentence examples

examples where the difference in evaluation is largest in either direction.

Tab. 3 shows two examples (cherry-picked from the top 10 sentence examples due to space limitations). We can see that in the first example, the PBMT-based system performs better on accurately translating a low-frequency named entity, while in the second example the NMT system accurately generates a multi-word expression with many frequent words. These concrete examples can both help reinforce our understanding of the patterns found in the holistic analysis above, or uncover new examples that may lead to new methods for holistic analysis.

3 Advanced Features

Here we discuss advanced features that allow for more sophisticated types of analysis using other sources of information than the references and system outputs themselves.

Label-wise Abstraction One feature that greatly improves the flexibility of analysis is `compare-mt`'s ability to do analysis over arbitrary word labels. For example, we can perform word accuracy analysis where we bucket the words by POS tags, as shown in 4. In the case of the PBMT vs. NMT analysis above, this uncovers the interesting fact that PBMT was better at generating base-form verbs, whereas NMT was better at generating conjugated verbs. This can also be applied to the n -gram analysis, finding which POS n -grams are generated well by one system or another, a type of analysis that was performed by Chiang et al. (2005) to understand differences in reordering between different systems.

Labels are provided by external files, where there is one label per word in the reference and system outputs, which means that generating these labels can be an arbitrary pre-processing step performed by the user without any direct modifications to the `compare-mt` code itself. These labels do not have to be POS tags, of course, and can

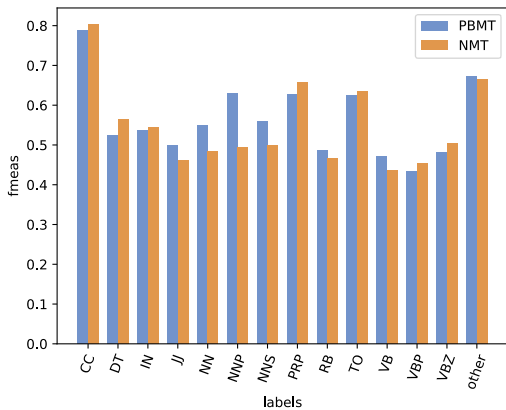


Figure 3: Word F-measure bucketed by POS tag.

also be used for other kinds of analysis. For example, one may perform analysis to find accuracy of generation of words with particular morphological tags (Popović et al., 2006), or words that appear in a sentiment lexicon (Mohammad et al., 2016).

Source-side Analysis While most analysis up until this point focused on whether a particular word on the *target* side is accurate or not, it is also of interest what source-side words are or are not accurately translated. `compare-mt` also supports word accuracy analysis for source-language words given the source language input file, and alignments between the input, and both the reference and the system outputs. Using alignments, `compare-mt` finds what words on the source side were generated correctly or incorrectly on the target side, and can do aggregate word accuracy analysis, either using word frequency or labels such as POS tags.

Word Likelihood Analysis Finally, as many recent methods can directly calculate a log likelihood for each word, we also provide another tool `compare-ll` that makes it possible to perform holistic analysis of these log likelihoods. First, the user creates a file where there is one log likelihood for each word in the reference file, and then, like the word accuracy analysis above, `compare-ll` can calculate aggregate statistics for this log likelihood based on word buckets.

Extending `compare-mt` One other useful feature is `compare-mt`'s ability to be easily extended to new types of analysis. For example,

- If a user is interested in using a different evaluation metric, they could implement a new instance of the `Scorer` class and use it for

both aggregate score analysis (with significance tests), sentence bucket analysis, or sentence example analysis.

- If a user wanted to bucket words according to a different type of statistic or feature, they could implement their own instance of a `Bucketer` class, and use this in the word accuracy analysis.

4 Example Use-cases

To emphasize `compare-mt`'s practical utility, we also provide examples of how it has *already* been used in analyses in published research papers: Figs. 4 and 5 of Wang et al. (2018) use sentence-level bucketed analysis. Tab. 7 of Qi et al. (2018) and Tab. 8 of Michel and Neubig (2018) show the results of *n*-gram analysis. Fig. 2 of Qi et al. (2018), Fig. 4 of Sachan and Neubig (2018), Tab. 5 of Kumar and Tsvetkov (2019) show the results of word accuracy analysis.

5 Conclusion

In this paper, we presented an open-source tool for holistic analysis of the results of machine translation or other language generation systems. It makes it possible to discover salient patterns that may help guide further analysis.

`compare-mt` is evolving, and we plan to add more functionality as it becomes necessary to further understand cutting-edge techniques for MT. One concrete future plan includes better integration with example-by-example analysis (after doing holistic analysis, clicking through to individual examples that highlight each trait), but many more improvements will be made as the need arises.

Acknowledgements: The authors thank the early users of `compare-mt` and anonymous reviewers for their feedback and suggestions (especially Reviewer 1, who found a mistake in a figure!). This work is sponsored in part by Defense Advanced Research Projects Agency Information Innovation Office (I2O) Program: Low Resource Languages for Emergent Incidents (LORELEI) under Contract No. HR0011-15-C0114. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- Koichi Akabe, Graham Neubig, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2014. Discriminative language models as a tool for machine translation error analysis. In *Proc. COLING*, pages 1124–1132.
- Wilker Aziz, Sheila Castilho, and Lucia Specia. 2012. Pet: a tool for post-editing and assessing machine translation. In *Proc. LREC*, pages 3982–3987.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. ICLR*.
- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. [Evaluating Discourse Phenomena in Neural Machine Translation](#). In *Proc. NAACL*, pages 1304–1313, New Orleans, Louisiana. Association for Computational Linguistics.
- Luisa Bentivogli, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. [Neural versus phrase-based machine translation quality: a case study](#). In *Proc. EMNLP*, pages 257–267, Austin, Texas. Association for Computational Linguistics.
- Alexandra Birch, Miles Osborne, and Phil Blunsom. 2010. Metrics for mt evaluation: evaluating reordering. *Machine Translation*, 24(1):15–26.
- Konstantinos Chatzitheodorou and Stamatis Chatzistamatis. 2013. COSTA MT evaluation tool: An open toolkit for human machine translation evaluation. *The Prague Bulletin of Mathematical Linguistics*, 100(1):83–89.
- David Chiang, Adam Lopez, Nitin Madnani, Christof Monz, Philip Resnik, and Michael Subotin. 2005. The hiero machine translation system: Extensions, evaluation, and analysis. In *Proc. EMNLP*, pages 779–786.
- Steve DeNeefe, Kevin Knight, and Hayward H. Chan. 2005. [Interactively exploring a machine translation model](#). In *Proc. ACL*, pages 97–100. Association for Computational Linguistics.
- Michael Denkowski and Alon Lavie. 2011. Meteor 1.3: Automatic Metric for Reliable Optimization and Evaluation of Machine Translation Systems. In *Proc. WMT*, pages 85–91.
- Yanzhuo Ding, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Visualizing and understanding neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1150–1159.
- Ahmed El Kholly and Nizar Habash. 2011. Automatic error analysis for morphologically rich languages. In *Proc. MT Summit*, pages 225–232.
- Christian Federmann. 2012. Appraise: an open-source toolkit for manual evaluation of mt output. *The Prague Bulletin of Mathematical Linguistics*, 98(1):25–35.
- Mark Fishel, Rico Sennrich, Maja Popović, and Ondřej Bojar. 2012. Terrorcat: a translation error categorization-based mt quality metric. In *Proc. WMT*, pages 64–70.
- Mary Flanagan. 1994. Error classification for mt evaluation. In *Proc. AMTA*, pages 65–72.
- Meritxell González, Jesús Giménez, and Lluís Màrquez. 2012. A graphical interface for mt evaluation and error analysis. In *Proceedings of the ACL 2012 System Demonstrations*, pages 139–144.
- Pierre Isabelle, Colin Cherry, and George Foster. 2017. [A challenge set approach to evaluating machine translation](#). In *Proc. EMNLP*, pages 2476–2486.
- Hideki Isozaki, Tsutomu Hirao, Kevin Duh, Katsuhito Sudoh, and Hajime Tsukada. 2010. Automatic evaluation of translation quality for distant language pairs. In *Proc. EMNLP*, pages 944–952.
- Ondřej Klejch, Eleftherios Avramidis, Aljoscha Burchardt, and Martin Popel. 2015. Mt-compareval: Graphical evaluation interface for machine translation development. *The Prague Bulletin of Mathematical Linguistics*, 104(1):63–74.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proc. EMNLP*, pages 388–395.
- Philipp Koehn. 2010. *Statistical Machine Translation*. Cambridge Press.
- Phillip Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proc. HLT*, pages 48–54.
- Sachin Kumar and Yulia Tsvetkov. 2019. [Von mises-fisher loss for training sequence to sequence models with continuous outputs](#). In *Proc. of ICLR*.
- Jaesong Lee, Joong-Hwi Shin, and Jun-Seok Kim. 2017. [Interactive visualization and manipulation of attention-based neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 121–126, Copenhagen, Denmark. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81.
- Adam Lopez and Philip Resnik. 2005. Pattern visualization for machine translation output. In *Proceedings of HLT/EMNLP 2005 Interactive Demonstrations*, pages 12–13.
- Nitin Madnani. 2011. ibleu: Interactively debugging and scoring statistical machine translation systems. In *2011 IEEE Fifth International Conference on Semantic Computing*, pages 213–214. IEEE.

- Inderjeet Mani. 1999. *Advances in automatic text summarization*. MIT press.
- Paul Michel and Graham Neubig. 2018. [MTNT: A testbed for machine translation of noisy text](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium.
- Margaret Mitchell, Jesse Dodge, Amit Goyal, Kota Yamaguchi, Karl Stratos, Xufeng Han, Alyssa Mensch, Alex Berg, Tamara Berg, and Hal Daume III. 2012. [Midge: Generating image descriptions from computer vision detections](#). In *Proc. EACL*, pages 747–756.
- Saif M Mohammad, Mohammad Salameh, and Svetlana Kiritchenko. 2016. How translation alters sentiment. *Journal of Artificial Intelligence Research*, 55:95–130.
- Mathias Müller, Annette Rios, Elena Voita, and Rico Sennrich. 2018. [A large-scale test set for the evaluation of context-aware pronoun translation in neural machine translation](#). In *Proc. WMT*, pages 61–72, Belgium, Brussels. Association for Computational Linguistics.
- Masaki Murata, Kiyotaka Uchimoto, Qing Ma, Toshiyuki Kanamaru, and Hitoshi Isahara. 2005. Analysis of machine translation systems’ errors in tense, aspect, and modality. In *Proc. PACLIC*.
- Sudip Kumar Naskar, Antonio Toral, Federico Gaspari, and Andy Way. 2011. A framework for diagnostic evaluation of mt based on linguistic checkpoints. *Proc. MT Summit*, pages 529–536.
- Graham Neubig and Junjie Hu. 2018. [Rapid adaptation of neural machine translation to new languages](#). In *Proc. EMNLP*, Brussels, Belgium.
- Alice H Oh and Alexander I Rudnicky. 2000. Stochastic language generation for spoken dialogue systems. In *Proceedings of the 2000 ANLP/NAACL Workshop on Conversational systems-Volume 3*, pages 27–32. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. ACL*, pages 311–318.
- Maja Popović. 2015. [chrF: character n-gram f-score for automatic mt evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Maja Popović, Adrià de Gispert, Deepa Gupta, Patrik Lambert, Hermann Ney, José B. Mariño, Marcello Federico, and Rafael Banchs. 2006. Morphosyntactic information for automatic error analysis of statistical machine translation output. In *Proc. WMT*, pages 1–6.
- Maja Popović and Hermann Ney. 2007. Word error rates: Decomposition over POS classes and applications for error analysis. In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 48–55.
- Maja Popović and Hermann Ney. 2011. Towards automatic error analysis of machine translation output. *Computational Linguistics*, 37(4):657–688.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. When and why are pre-trained word embeddings useful for neural machine translation? In *Proc. NAACL*, New Orleans, USA.
- Ehud Reiter and Robert Dale. 2000. *Building natural language generation systems*. Cambridge university press.
- Devendra Sachan and Graham Neubig. 2018. [Parameter sharing methods for multilingual self-attentional translation models](#). In *Proc. WMT*, Brussels, Belgium.
- Rico Sennrich. 2017. [How Grammatical is Character-level Neural Machine Translation? Assessing MT Quality with Contrastive Translation Pairs](#). In *Proc. EACL*, pages 376–382, Valencia, Spain.
- Sara Stymne. 2011. BLAST: A tool for error analysis of machine translation output. In *Proceedings of the ACL-HLT 2011 System Demonstrations*, pages 56–61.
- David Vilar, Jia Xu, Luis Fernando d’Haro, and Hermann Ney. 2006. Error analysis of statistical machine translation output. In *Proc. LREC*, pages 697–702.
- Xinyi Wang, Hieu Pham, Pengcheng Yin, and Graham Neubig. 2018. [A tree-based decoder for neural machine translation](#). In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium.
- Jonathan Weese and Chris Callison-Burch. 2010. Visualizing data structures in parsing-based machine translation. *The Prague Bulletin of Mathematical Linguistics*, 93:127–136.
- Daniel Zeman, Mark Fishel, Jan Berka, and Ondřej Bojar. 2011. Addicter: What is wrong with my translations? *The Prague Bulletin of Mathematical Linguistics*, 96(1):79–88.
- Ming Zhou, Bo Wang, Shujie Liu, Mu Li, Dongdong Zhang, and Tiejun Zhao. 2008. [Diagnostic evaluation of machine translation systems using automatically constructed linguistic check-points](#). In *Proc. COLING*, pages 1121–1128, Manchester, UK. Coling 2008 Organizing Committee.